

COMMISSIONING AND FIRST PERFORMANCE OF THE LHC BEAM INSTRUMENTATION

R. Jones, CERN, Geneva, Switzerland.

Abstract

This paper will outline the progress with LHC commissioning to date, detailing the performance achieved with all the main LHC beam instrumentation systems. It will include an overview of the beam loss system and its role in machine protection, along with that of the beam position system and its use for automatic orbit control. Results will be shown from the highly sensitive base band tune system as well as the bunch-by-bunch and DC beam current transformer systems, the synchrotron light monitoring systems, the wire scanner system and OTR screens.

INTRODUCTION

The first beams circulated in the LHC on 10th September 2008 in full view of the world's media. This was a resounding success, and was followed a few days later with RF captured beams circulating for several hours. Nine days on, a poor superconducting splice overheated during a hardware test at high current, creating an arc that pierced the helium containment vessel, with severe consequences. 14 months later, after a major magnet repair and consolidation programme, the LHC was once again cold and ready to take beam. Despite these efforts and the installation of a completely new Quench Protection System (QPS) its energy is currently limited to 3.5 TeV per beam. Further consolidation, requiring a long shutdown in 2012, will be necessary before it can run at its design energy of 7 TeV per beam.

A one month run at the end of 2009 saw the LHC quickly advance with optics, collimation and working point studies at its 450 GeV injection energy. Ramp commissioning to 1.18 TeV followed, ending in collisions at 1.18 TeV per beam in all of its four main experiments.

Further consolidation work was carried out on both the machine and experiments for the first two months of 2010, before the start of a 2 year run at the end of February. To date (late April 2010) the LHC is routinely ramping 2×10^{10} protons to 3.5 TeV and colliding with a β^* of 2 m in all four experiments. The emphasis now is on increasing the total intensity in the machine.

This rapid progress with beam commissioning is in a large part down to the very good beam instrumentation with which the LHC is equipped. These systems and their contributions to this commissioning effort will be the subject of the remainder of this paper.

THE LHC BTV SYSTEM

There are a total of 37 TV beam observation systems (BTV) of 7 different types installed in the LHC. Each BTV station is equipped with two screens; one a 1 mm thick alumina plate (scintillator) and the second a 12 μ m

thick titanium foil (to produce Optical Transition Radiation or OTR). The alumina plates are very sensitive and can observe single bunches of well below 10^9 particles, but due to their thickness significantly perturb the beam. The number of photons emitted by the OTR is much less than that of the alumina screen; on the other hand the perturbation to the beam is minimal, allowing multiple monitors to be used at the same time, as well as multi-turn observation.

First Results from the LHC BTV System

During initial commissioning the beam was steered through the transfer lines and the different sections of the machine using the alumina screens, most of the time producing very clear but completely saturated images (see Figure 1).




Figure 1 : The first full turn in the LHC as seen by the BTV system (10/9/2008).

After this first step the OTR screens replaced the alumina screens to reduce the blow-up of the beam, reduce the radiation due to beam losses, and produce images well suited for analysis with good linearity and no saturation. There was a possibility of observing OTR emission for even the lowest intensity pilot bunches (2×10^9 protons) due to the high sensitivity of the standard CCD cameras.

All the BTV monitors have performed very well and were extensively used during the synchronization tests and for the first beams in the LHC. The large 1 m diameter alumina screen in the dump line is still continually used to verify the correct functioning of the beam dump system (see Figure 2).

Due to the expected radiation levels in the injection regions, a gradual replacement of the standard CCD cameras with less-sensitive, radiation-hard cameras is foreseen. This will result in the loss of ability to observe pilot bunches in OTR mode.




Figure 2: Checking the LHC beam dump dilution kickers with the dump line alumina screen. Predicted and measured positions of the 10 bunches used are shown.

THE LHC BLM SYSTEM

The purpose of the LHC BLM system is threefold: to protect the LHC equipment from damage; to avoid beam induced magnet quenches; and to observe losses for machine parameter tuning such as the adjustment of collimators. The protection requirements led to the placement of detectors at all likely loss locations, resulting in a total of some 4000 installed monitors. The majority of these are located around the quadrupole magnets and in the collimator regions. The signals of almost all monitors are compared with pre-defined threshold values which, if exceeded, result in a retraction of the beam permit signal and consequently a beam dump.

The LHC BLM Acquisition System

The majority of the LHC BLMs are 50 cm long, 1.5 litre, nitrogen filled ionisation chambers. These have been optimized to give an ion collection time of 85 μ s, i.e., less than one LHC turn. They are located around each quadrupole magnet (six per quadrupole), in the collimator regions and at other aperture restrictions around the machine. In order to fulfil the three criteria mentioned above, the system has to cope with a total dynamic range of some 10^{13} . This is achieved by combining the ionisation chambers with secondary emission monitors (SEMs) having $\sim 30000\times$ smaller gain [1].

The same acquisition system is used for both the ionisation chambers and SEMs, and is based on current to frequency conversion. The signal from the detector is integrated and compared to a predefined threshold. Once this threshold is reached a one-shot is fired, which switches on a constant current source to discharge the integrating capacitor until a lower threshold is attained and the current source is switched off. In this way the frequency of the one-shot reset output is proportional to

the input current from the monitor. The frequency ranges from 0.01 Hz to 1 MHz for currents of 3 pA to 300 μ A respectively with linearity better than 5%. The analogue front-end counts these reset pulses using an 8-bit asynchronous counter, with each count representing a fixed amount of charge. Every 40 μ s, the count from all six channels for a given quadrupole is multiplexed into a Manchester Encoded serial stream and transmitted to the surface using a Gigabit Optical Link with cyclic redundancy. The surface electronics calculates the integrated loss values for time periods of between 80 μ s and 100 s and compares them to a table of threshold values that depend on the loss duration and beam energy.

The beam loss monitor acquisition is an integral part of the machine protection system, and for losses occurring on a time scale of less than 10 ms it is the only loss detection system available for the LHC. For this reason the failure rate and availability requirements are very stringent and have been evaluated using the Safety Integrity Level (SIL) approach. The system has been calculated to reach the SIL3 level, corresponding to a 10^{-3} per year probability of failing to detect a dangerous beam loss. This is achieved by duplicating the signal treatment chain for all elements after the current-to-frequency conversion, incorporating error correction and detection techniques, and constantly monitoring the availability of all monitors.

LHC BLM System Performance

There were two beam-induced triggers of the quench protection system during the injection tests of 2008, which allowed an attempt at quench reconstruction using the BLM system. The loss pattern observed in each case was almost in the centre of a bending magnet in an area with sufficient BLM coverage. Knowing the bunch intensity, impact location and loss distribution widths, it was possible to compare the measured results with GEANT4 simulations (see Figure 3). This showed a discrepancy factor of ~ 1.5 , as a result of which the threshold values for quench prevention have been raised by $\sim 50\%$.




Figure 3: Comparison of measured and simulated shower evolution (data points indicate monitor locations).

The LHC BLM system is the main tool for setting-up the collimation system, which is essential for protecting the machine against damage. At the high energy and relatively low emittances of the LHC, this damage limit is quickly reached with only a few pilot bunches of 5×10^9 circulating in the machine.

Figure 4 shows the measured beam loss from all monitors for an unstable beam. The large dynamic range of the BLM system is clearly visible, with the results showing that the collimation system is performing as designed. Losses are localised to the collimator regions, giving a cleaning efficiency of better than 99.98%.

The fluctuation in the noise floor of the BLMs comes from the different cable lengths used to transport the signal from the chamber to the front-end electronics, which can be up to 700 m away in some cases. This is still below the quench level for beams at energies up to 5 TeV, but may have to be improved for nominal running at 7 TeV. One option being investigated is the use of a radiation-hard application-specific integrated circuit (ASIC) to replace the existing current-to-frequency converter, allowing electronic installation much closer to the monitor and so significantly reducing the cable lengths required.




Figure 4: Beam losses throughout the LHC ring as measured during a test of the collimation system.

THE LHC BPM SYSTEM

The LHC BPM system is comprised of 1054 beam position monitors, the majority of which (912) are 24mm button electrode BPMs located in all arc quadrupole cryostats (SSS). The remaining BPMs are enlarged (34mm or 40mm) button electrode BPMs mainly for the stand-alone quadrupoles, or stripline electrode BPMs used either for their directivity in the common beam pipe regions or for their higher signal level in the large diameter vacuum chambers around the dump lines.

The beam position acquisition electronics is split into two parts, an auto-triggered, analogue, position to time normaliser, [2] which sits in the tunnel; and an integrator/digitiser/processor VME module on the surface. Each BPM measures in both the horizontal and vertical plane, resulting in a total of 2156 electronic channels [3].

First Results from the LHC BPM System

The beam threading mode, also known as FIFO mode, was used to provide the first turn position data to the operator interface. This is a totally asynchronous acquisition mode, where any triggers obtained within a specified gate are stored, processed & published. From the very first shot into the LHC, the BPM system gave excellent results while operating in this mode. Combined with powerful orbit correction software it allowed quick diagnostics of BPM polarity and machine optics errors.

Test of the Intensity Limit for the BPM System

The LHC BPM system can work in one of 2 sensitivity ranges – high or low. The threshold for the sensitivity change is conceived such that any reflections at the level of -20dB do not retrigger the system. In the high sensitivity setting the BPM system was foreseen to work correctly with bunch intensities in the range 2×10^9 to 2×10^{10} protons, with the low sensitivity setting working from 2×10^{10} protons upwards. In order to test the lower limit for auto-trigger detection the bunch intensity was varied from below 1×10^9 to around 3×10^9 , while counting the number of correctly triggered BPMs in the arc of Sector 2-3. The lower limit for correct BPM functioning was found to be $\sim 1.5 \times 10^9$ protons per bunch.

Early Performance of the LHC BPM System

The performance of the LHC BPM system has been characterised for low intensity, single bunch operation, both in trajectory and in orbit mode. From the trajectory analysis of 31 injections during the synchronisation tests in 2008 the overall rms variation in position over a timescale of some ten minutes was generally seen to be between 150-400 μm . This order of magnitude is consistent with the electronic noise estimations for bunch intensities in the range 2×10^9 to 5×10^9 protons.




Figure 5: Orbit stability (12/9/2008).

Once the beam started to circulate, the asynchronous orbit acquisition of the BPM system (IIR mode) could be used. This provides an update of the average orbit at 25 Hz. Figure 5 shows the results of some 15 minutes worth of horizontal orbit data, with similar results obtained in the vertical plane. It can be seen that the rms noise on the measurement for the majority of the machine is between 5 and 15 μm . The alternating high/low noise peaks follow the beta function (with 45° BPM sampling), indicating that a large fraction of this noise results from the beam itself, with an additional contribution from the known COD power converter current ripple. The resolution of the BPM system in orbit mode with a single pilot bunch is therefore better than $\sim 5 \mu\text{m}$.

An oscillation of the position in Point 2 and a drift of the position in Point 4 are also clearly visible. These variations were traced to changes in the air temperature of the surface electronics racks in these locations and may be unacceptable for nominal LHC operation. Investigations are therefore underway to try to eliminate this problem.

Commissioning of Orbit Feedback

Full commissioning and characterisation of the LHC orbit feedback system has recently started. The system currently has a closed loop bandwidth of 0.1 Hz and is continually supplied with orbit data from all 1054 BPMs at 25 Hz. A comparison of the orbit drifts observed with and without feedback during energy ramps to 3.5 TeV is shown in Figure 6.




Figure 6: Comparison of orbit stability and drift with and without orbit feedback

It can be seen that although the rms orbit is significantly improved with the feedback on, there are still orbit perturbations that remain uncorrected. Analysis has shown that these are due to localised distortions that are not fully taken into account with the global correction currently being applied. Work is ongoing to address this problem and to integrate orbit feedback into day-to-day operation.

THE LHC TUNE AND CHROMATICITY MEASUREMENT SYSTEMS

The LHC base-line tune, chromaticity and coupling measurement system consists of three independent acquisition chains per beam and relies on the diode-based, base-band-tune (BBQ) technique [4] developed for the LHC and now also used in all CERN synchrotrons. The layout of these systems is described in detail in [5].

Of the three systems per beam, the first is dedicated to passive beam spectrum observation and ensures a continuous data logging for post-mortem analysis, passive beam quality monitoring and fixed displays in the control room. This passive system observes any residual beam oscillations or externally produced beam excitation.

The second system is intended for “on-demand” tune measurement, with excitation provided using dedicated tune kickers or using fast frequency sweeps (“chirp” signals) via the LHC transverse damper.

The third system is dedicated to PLL tune operation, again using the transverse damper as excitation source. While initial PLL testing has taken place, this system remains to be fully commissioned.

First Results from the LHC Tune System

Fig. 7 shows a typical tune spectrum captured at injection energy with no externally applied excitation. The calibration of the signal amplitude was performed by applying a known single frequency tone via a stripline kicker. The remarkable sensitivity of the BBQ system is clearly visible, with oscillations visible down to the tens-of-nanometres level.




Figure 7: Typical tune spectrum resulting from residual beam oscillations. Also visible is the unknown perturbation source known as “the hump”.

The figure also shows a perturbation of unknown origin in the vertical plane known as “The Hump”. This wandering excitation has a quickly varying component, showing up as a wide frequency hump in the spectrum, and a slowly varying component with a ~ 7 minute period

which moves the whole distribution around. This causes problems whenever it crosses the tune line, resulting in emittance blow-up and beam-loss seen through a reduction in beam lifetime. Discovering the source of this is excitation is therefore becoming a priority.

Commissioning Tune Feedback

With full pre-cycling, the fill-to-fill tune stability during ramping is now typically $2\text{--}3 \times 10^{-3}$. However, during early commissioning the variations were up to 0.1 and are still frequently up to 0.02 due to partial magnet pre-cycles after access or sector trips for example. Tune feedback was therefore crucial early on and is still routinely used for ramping.

With the phased-locked loop (PLL) system not yet fully commissioned, the feedback relies on a peak fit of the BBQ FFT spectra to provide the tune input. The closed loop bandwidth is currently between 0.1-0.3 Hz and limited by the interaction of real-time corrections applied to the trim quadrupoles and their new QPS system. Figure 8 shows examples of ramps performed with and without tune feedback. With the feedback on a peak-to-peak tune stability of 10^{-3} has been achieved.




Figure 8: Tune spectra during an energy ramp without tune feedback (top) and with tune feedback (bottom).

Chromaticity Measurement

Standard chromaticity measurement performed at fixed energy by measuring the tune changes observed via the BBQ system while varying the RF frequency was operational from day 1.

The continuous measurement of chromaticity during the ramp is performed by demodulating the tune signals observed while applying sinusoidal frequency trims (see Figure 9). The modulation amplitude has had to be increased from $\Delta p/p = 10^{-5}$ at 2.5 Hz to $\Delta p/p = 10^{-4}$ at 2 Hz to overcome the tune jitter at injection, where the residual tune stability is not better than 3×10^{-4} . Using this technique the nominal chromaticity resolution of ± 2 units has been reached, with the results fed forward for the next ramps. Chromaticity feedback has not yet been attempted.




Figure 9: Continuous chromaticity measurement during the snap-back at the start of the ramp.

THE LHC WIRESCANNER SYSTEMS

Eight linear wire scanners are installed in the LHC for measuring the transverse density profile of the beam. Each beam is equipped with two wire scanners in both horizontal and vertical planes, with one available for operation and the other intended as a fully functional back-up. The scanning speed of these systems is 1 ms^{-1} . Acquisition is possible in two different modes: the standard mode, using a logarithmically amplified PMT signal, to give the average beam size over all bunches; and the bunch-by-bunch mode which relies on a fast integrator ASIC developed for the LHCb experiment working at 40 MHz to integrate the PMT signal of each bunch individually and hence give its size. To date only the standard system has been commissioned with beam.

First results from the Wire Scanner System

The wire scanners have been extensively used from an early stage to monitor the beam size and calculate its emittance. Clean signals are obtained even with single, very low intensity pilot bunches of around 5×10^9 protons varying in size (sigma) between a few mm and $\sim 300 \mu\text{m}$. The resolution in the measurement of the rms beam size

under these conditions has not yet been fully evaluated, but is at the tens-of-microns level, with a reproducibility (in/out scans and repeated scans) of the same order.




Figure 10: In and out wire scans for the horizontal plane of Beam 2 performed at 3.5 TeV (26/4/2010).

The linearity of the wirescanner system was tested by comparing the centre of the distribution with that measured with the LHC orbit system during closed bump scans. Initial results show that this linearity is better than 5% for all scanners.

The limitations of the system have also been extensively studied [6]. Simulations corroborated by SPS measurements show that a 30 μm carbon wire scanned at 1ms^{-1} will survive (reaching a temperature less than 4000 K) with 25% of nominal LHC beam at the 450 GeV injection energy and 7% of nominal LHC beam at 7 TeV. These limits are a function of beam size, not energy. At high energy, however, this is not the limiting factor. Simulations have shown that, between 0.4% to 2% of nominal LHC beam at 7 TeV, the first downstream magnet reaches its quench limit during a scan. The wire scanner can therefore only be used to cross calibrate the other emittance measurement devices (synchrotron light monitor for protons and rest gas ionisation monitor for ions) with limited intensity in the machine.

THE LHC SYNCHROTRON LIGHT SYSTEMS

Synchrotron light is used by two systems in the LHC:

- The abort gap monitor, which ensures that the level of protons in the 3 μs abort gap, required to accommodate the rise time of the dump line extraction kickers, remains below the magnet quench level.
- The synchrotron light monitors, which provide a continuous image of the beam.

The light detected in both of these systems comes from multiple sources. A superconducting, 5 tesla, 2 period undulator is used to produce enough photons for imaging

single pilot bunches of 5×10^9 protons from the injection energy of 450 GeV to around 1.2 TeV. Above this energy the undulator radiation is shifted in wavelength into the ultraviolet and can no longer be imaged. From this stage, the dipole edge radiation from a neighbouring separation dipole takes over, and is used up to about 3 TeV. Above this energy, central dipole radiation from the same magnet is used.

The system therefore has to contend with not only the spectral and intensity variation with the increase in energy, but also the change in focus of the source.

Issues during the initial installation in 2008 led to a complete redesign of the system in 2009. This new system is built on a 4.8m long optical table located under the beam pipe and surrounded by heavy shielding to minimise future radiation problems. A motorized delay line is used to change the focus from the undulator to the dipole, while horizontal slits select the appropriate source point and optical density filter wheels control the light intensity. The table also contains a 28 m long calibration line, (the length being the distance from the beam light source to the imaging plane), used for image quality characterisation.

The LHC Abort Gap Monitoring System

The abort gap monitor is based on a gated photomultiplier (PMT), which receives $\sim 15\%$ of the total light collected. Its purpose is to detect the light generated by the relatively few protons which could cause a quench should they be kicked out during the beam dump kicker rise time. In order not to saturate during the passage of bunches in their nominal position, the PMT is gated off except during the 3 μs abort gap where the measurements are to be carried out.




Figure 11: A single pilot bunch of 5×10^9 debunching into the abort gap when the RF is turned off.

The PMT signal is integrated over thirty 100-ns bins throughout the gap using a fast integrator ASIC developed for the LHCb experiment and also used by the fast BCT, wirescanner and collision rate monitors. The population in each of these 100ns bins is averaged over 100 ms, giving a warning to the operators once this value exceeds 10% of the quench threshold. For protons near

1.2 TeV and ions at injection energy, an averaging time of 1 second will probably be necessary to cope with the low photon flux at these energies.

Figure 11 shows the results of a test where the RF was switched off with a single pilot bunch of 5×10^9 in the machine. Debunching of the beam is clearly visible, with the abort gap quickly being populated.

The LHC Synchrotron Light Monitor

The synchrotron light monitors are continuously used to monitor the beam size in the LHC. Initially, with low intensity beams, they were also the only instruments capable of showing that there was still something in the machine if the RF tripped and debunched the beam.

Figure 12 shows a comparison of the beam size as measured continuously by the synchrotron light monitor and periodically by the wire scanner during an LHC energy ramp to 3.5 TeV. There is good agreement between the measurements during the injection plateau and the early part of the ramp, where the focus is on the undulator radiation. As the energy rises, the increase in edge radiation from the separation dipole blurs the image, leading to a measured beam size larger than that given by the wire scanners. Adjusting the focussing trombone recovers some of the image quality, which is further improved by addition of a blue filter. Further fine tuning will be necessary to reach the 3-5% error level predicted by calculations for the vertical plane.

Using the synchrotron light monitor, it has been possible to study the emittance growth observed in the vertical plane of both beams (see Figure 12 before the start of the ramp). The source of this emittance growth has not yet been determined.




Figure 12 : A continuous measurement of the Beam 1 vertical size throughout the LHC energy ramp to 3.5 TeV. Wire scanner measurements are seen as squares.

THE LHC BCT SYSTEMS

Beam current transformers of two different kinds provide intensity measurements for the beams circulating in the LHC rings, as well as for the transfer lines from the SPS to LHC and from LHC to the dumps.

The LHC BCTDC System

The DC current transformers (BCTDC) are based on the principle of magnetic amplifiers and measure the mean current of the circulating beam [7]. Because of their operational importance, two independent systems are installed per ring. Four ranges, provided simultaneously, are used to cover the entire dynamic range of the beam from a few 10^9 to 5×10^{14} protons ($\sim 3 \mu\text{A}$ to $\sim 900 \text{ mA}$).

First results from the LHC BCTDC System

The main application used to view the BCTDC acquisition is a fixed display refreshed once per second. Figure 13 shows such an acquisition for Beam 2 on September 12th 2008, with 4 successive injections of a single pilot bunch ranging in intensity from 3.5×10^9 to 6×10^9 protons and circulation times of 100 seconds to 40 minutes.




Figure 13: First circulating beam in ring 2 as seen by the BCTDC (12/9/2008).

The noise on the signal corresponds to an rms of 7×10^8 protons for a one-second integration time, which is equivalent to $1.3 \mu\text{A}$. A similar value is found for the slow fluctuations observed over several hours. This result is very good and meets expectations, although it is also a reminder that the BCTDC can only measure single pilot bunches with a limited accuracy of around 20%. A small, negative offset of $\sim 2.5 \times 10^9$ ($4.5 \mu\text{A}$) was also observed at that time and is now automatically corrected for.

The LHC Fast BCT (BCTFR/BCTFD) Systems

The LHC fast BCT systems are capable of measuring bunch to bunch intensity on a turn by turn basis. Four such systems are installed in the LHC ring (BCTFR) – one operational system and a hot spare for each beam. Each system has two parallel measurement bandwidths (providing bunch to bunch and average turn intensity) acquiring simultaneously with two different gains. Two systems per beam are also installed in the LHC dump lines (BCTFD) to measure the total dumped beam intensity.

First results from the BCTFR System

All the fast BCTs have been commissioned in each of the abovementioned measurement modes [8]. The behaviour of the system with beam is already well understood, and the calibration agrees very well with what is obtained with the DCCT over the intensity range that has so far been used in the LHC (see Fig. 14).

The main issue that still needs to be addressed is the baseline restoration, which compensates for any droop and ensures that any offset is correctly taken into account. This is very important when producing the sum of intensities over a given turn, as any slight offset is magnified by 3564, corresponding to the number of 25-ns integration periods within the 89- μ s LHC revolution time. A new baseline restoration algorithm is currently being tested which should significantly improve this performance and make the system much less sensitive to the validity threshold settings that currently have to be used to ensure that only valid bunch signals are summed.

Noise levels at around 10^8 charges on the average intensity data produced every 20 ms have allowed lifetime measurements of several hundred hours to be calculated within several minutes. This has been very useful to tune-up and optimise the machine.

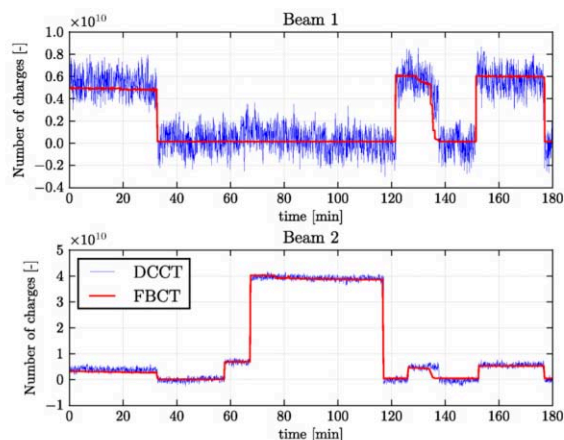


Figure 14: A comparison of fast BCT and DCCT data for various beam intensities.

SUMMARY

The commissioning **with beam** of the LHC has gone remarkably smoothly, from the first bunch injected in 2008 to the recent 30 hour physics fill at 3.5 TeV with a β^* of 2 m in all experiments. The excellent quality of the LHC beam instrumentation has had a major role to play in this success, with all major systems available with near full functionality from day 1. For the machine the emphasis will now be on providing ever more intense

beams while ensuring the protection of the machine. From the instrumentation side the work will switch from commissioning to optimisation.

A few remaining systems have yet to be fully commissioned, namely the LHC Schottky monitors and the LHC fast ionisation chambers for collision rate measurement in the two high-luminosity interaction regions. While first signals have been detected on both of these systems, they await higher intensities before the final commissioning can be completed.

ACKNOWLEDEMENTS

I would like to acknowledge all the members of the CERN Beam Instrumentation Group, both past and present, who have worked so hard to make these systems perform as well as they have. In the same manner I also thank all our external collaborators, who are too numerous to name, but without whom we would not have had such a successful LHC start-up with beam.

REFERENCES

- [1] B. Dehning et al., "The LHC Beam Loss Measurement System", 22nd Particle Accelerator Conference, Albuquerque, NM, USA, 25 - 29 Jun 2007, p. 4192.
- [2] D. Cocq, "The Wide Band Normaliser : a New Circuit to Measure Transverse Bunch Position in Accelerators and Colliders", Nucl. Instrum. Methods Phys. Res., A **416**, 1 (1998) 1.
- [3] R. Jones et al., "The LHC Orbit and Trajectory System", 6th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators, Mainz, Germany, 5-7 May 2003, p. 187.
- [4] M. Gasior and R. Jones, "High Sensitivity Tune Measurement by Direct Diode Detection", 7th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators, Lyon, France, 6-8 Jun 2005, p. 310.
- [5] M. Gasior et al., "An Overview of the LHC Transverse Diagnostics Systems", CERN-LHC-PROJECT-Report-1166 (2008).
- [6] M. Sapinski and T. Kroyer, C, 13th Beam Instrumentation Workshop, Tahoe City, CA, USA, 04 -08 May 2008, pp. 383-388.
- [7] P. Odier et al., "The DCCT for the LHC Beam Intensity Measurement", 9th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators, Basel, Switzerland, 25 - 27 May 2009.
- [8] D. Belohrad et al., "The LHC Fast BCT System – A Comparison of Design Parameters with Initial Performance", these proceedings.